



# Deep Learning for Whole Slide Image Analysis: An Overview

Neofytos Dimitriou<sup>1\*</sup>, Ognjen Arandjelović<sup>1</sup> and Peter D. Caie<sup>2</sup>

<sup>1</sup> School of Computer Science, University of St Andrews, St Andrews, United Kingdom, <sup>2</sup> School of Medicine, University of St Andrews, St Andrews, United Kingdom

The widespread adoption of whole slide imaging has increased the demand for effective and efficient gigapixel image analysis. Deep learning is at the forefront of computer vision, showcasing significant improvements over previous methodologies on visual understanding. However, whole slide images have billions of pixels and suffer from high morphological heterogeneity as well as from different types of artifacts. Collectively, these impede the conventional use of deep learning. For the clinical translation of deep learning solutions to become a reality, these challenges need to be addressed. In this paper, we review work on the interdisciplinary attempt of training deep neural networks using whole slide images, and highlight the different ideas underlying these methodologies.

**Keywords:** digital pathology, computer vision, oncology, cancer, machine learning, personalized pathology, image analysis

## OPEN ACCESS

### Edited by:

Inti Zlobec,  
University of Bern, Switzerland

### Reviewed by:

Pier Paolo Piccaluga,  
University of Bologna, Italy  
Thomas Menter,  
University Hospital of Basel,  
Switzerland

### \*Correspondence:

Neofytos Dimitriou  
neofytosd@gmail.com

### Specialty section:

This article was submitted to  
Pathology,  
a section of the journal  
Frontiers in Medicine

**Received:** 07 June 2019

**Accepted:** 29 October 2019

**Published:** 22 November 2019

### Citation:

Dimitriou N, Arandjelović O and  
Caie PD (2019) Deep Learning for  
Whole Slide Image Analysis: An  
Overview. *Front. Med.* 6:264.  
doi: 10.3389/fmed.2019.00264

## 1. INTRODUCTION

The adoption of digital pathology into the clinic will arguably be one of the most disruptive technologies introduced into the routine working environment of pathologists. Digital pathology has emerged with the digitization of patient tissue samples and in particular the use of digital whole slide images (WSIs). These can be distributed globally for diagnostic, teaching, and research purposes. Validation studies have shown correlation between digital diagnosis and glass based diagnosis (1, 2). However, although multiple whole slide scanners are currently available on the market, to date only Philips Ultra-fast scanner has been approved by regulatory bodies for the use in primary diagnosis (3).

Recently, NHS Greater Glasgow and Clyde, one of the largest single site pathology services in Europe, has begun proceedings to undergo full digitization. As the adoption of digital pathology becomes wider, automated image analysis of tissue morphology has the potential to further establish itself in pathology and ultimately decrease the workload of pathologists, reduce turnaround times for reporting, and standardize clinical practices. For example, known or novel biomarkers and histopathological features can be automatically quantified (4–12). Furthermore, deep learning techniques can be employed to recognize morphological patterns within the specimen for diagnostic and triaging purposes (13–16).

Successful application of deep learning to WSIs has the potential to create new clinical tools that surpass current clinical approaches in terms of accuracy, reproducibility, and objectivity while also providing new insights on various pathologies. However, WSIs are multi-gigabyte images with typical resolutions of 100,000 × 100,000 pixels, present high morphological variance, and often contain various types of artifacts. These conditions preclude the direct application of conventional deep learning techniques. Instead, practitioners are faced with two non-trivial challenges. On the one hand, the visual understanding of the images, impeded by the morphological variance, artifacts, and typically small data sets, and, on the other hand, the inability of the current state of the

hardware to facilitate learning from images with such high resolution, thereby requiring some form of dimensionality reduction to the images. These two problems are sometimes referred to as the *what* and *where* problems (17, 18). In this paper, we first discuss important aspects and challenges of WSIs, and then delve deeper into the different approaches to the two aforementioned problems.

## 2. WHOLE SLIDE IMAGES

### 2.1. Tissue Visualization

The majority of WSIs are captured using brightfield illumination, such as for slides stained with clinically routine haematoxylin and eosin (H&E). The wider accessibility of H&E stained WSIs, compared to more bespoke labeling reagents, at present makes this modality more attractive for deep learning applications. H&E stained tissue is excellent for the characterization of morphology within a tissue sample which corroborates to its long use in clinical practice.

However, H&E stained slides lack *in situ* molecular data associated with a cell. In contrast, this is possible with protein visualization through immunolabeling. The labeling of multiple cell types and their protein expression can be observed with multiplexed immunofluorescence (IF) which provides valuable information in cancer research and particularly in immuno-oncology (7, 10, 19). Nevertheless, the analysis of IF labeled slides using deep learning techniques is impeded by the limited availability of IF WSI data sets. Potential causes of the data scarcity include the expense of reagents and of access to fluorescence scanners, as well as the enormous IF WSI size which can sometimes exceed 10 gigabytes per image.

### 2.2. Data Availability

Unlike in numerous other fields which have adopted supervised deep learning techniques (20, 21), labeled data is more difficult to obtain in digital pathology, thereby challenging the practicability of supervised approaches. Despite wider data publication in the recent years (13, 22–26), much of the published work still employs proprietary WSI data sets (27).

### 2.3. Image Format

There are currently multiple whole slide scanners from different vendors available on the market with the capacity for both brightfield and fluorescence imaging. Each scanner captures images using different compression types and sizes, illumination, objectives, and resolution and also outputs the images in a different proprietary file format. The lack of a universal image format can delay the curation of large data sets. The field of radiology has overcome this issue with the adoption of DICOM open source file formats allowing large image data sets to be accessed and interrogated (28, 29). Digital pathology is yet to widely adopt a single open source file format although work and discussions are continually progressing toward this end (30, 31).

### 2.4. Artifacts and Color Variability

To be clinically translatable, deep learning algorithms must work across large patient populations and generalize over image

artifacts and color variability in staining (32). Artifacts can be introduced throughout the entire sample preparation workflow as well as during the imaging process. These can include ischemia times, fixation times, microtome artifact, staining reagent variability as well as imaging artifacts from uneven illumination, focusing, image tiling and fluorescence deposits and bleed-through. Examples of such artifacts are shown in **Figure 1**.

Through training, the human brain can become adept at ignoring artifacts and staining variability, and honing in the visual information necessary for an accurate diagnosis. To facilitate an analogous outcome in deep learning models, there are generally two approaches that can be followed. The first involves explicit removal of artifacts (e.g., using image filters), as well the normalization of color variability (33). In contrast, the second approach takes on a less direct strategy, augmenting data with often synthetically generated data which captures a representative variability in artifacts and staining, making their learning an integral part of the training process. Both approaches have been employed with some success to correct the variation from batch effect or from archived clinical samples from different clinics (34) though this finding has not been universal (15).

## 3. DEEP LEARNING

### 3.1. Patch Extraction

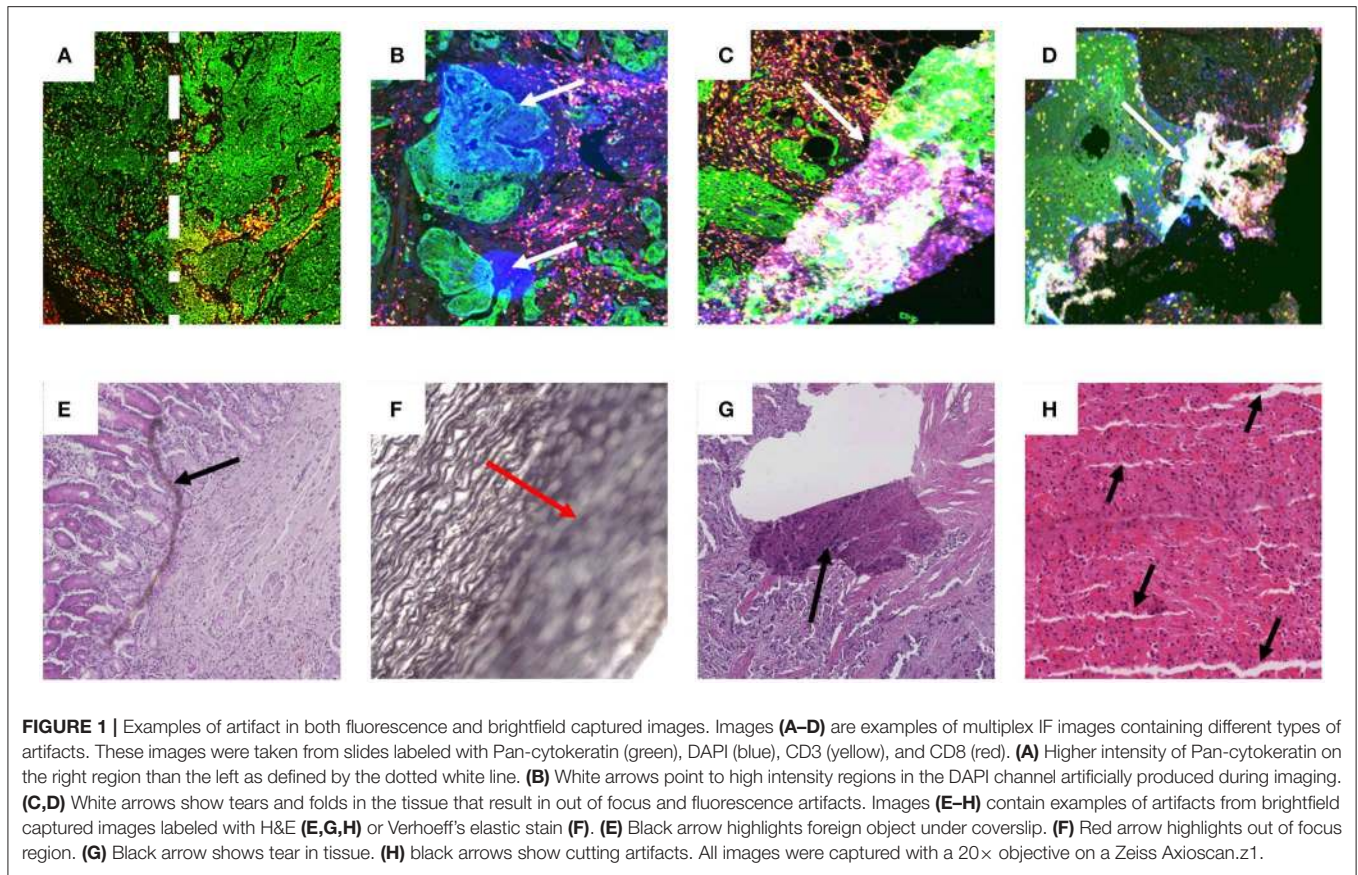
Most successful approaches to training deep learning models on WSIs do not use the whole image as input and instead extract and use only a small number of patches (6, 23–25). Image patches are usually square regions with dimensions ranging from  $32 \times 32$  pixels up to  $10,000 \times 10,000$  pixels with the majority of approaches using image patches of around  $256 \times 256$  pixels (6, 25, 35). This approach to reducing the high dimensionality of WSIs can be seen as human guided feature selection. The way patches are selected constitutes one of the key areas of research for WSI analysis. Existing approaches can be grouped based on whether they employ annotations and at which level (see **Figure 2**).

#### 3.1.1. Patch Level Annotation

Patch level annotations enable strong supervision since all of the extracted training patches have class labels. Typically, patch based annotations are derived from pixel level annotations which requires experts to annotate all pixels. For instance, given a WSI which contains cancerous tissue, a pathologist would need to localize and annotate all cancerous cells.

A simple approach to patch based learning would make use of all tiled (i.e., non-overlapping) patches. Nevertheless, this simplicity comes at the cost of excessive computational and memory overhead, along with a number of other issues, such as imbalanced classes and slow training. Randomly sampling patches may lead to an even higher class imbalance considering how in most cases a patch is much smaller than the original WSI. It is therefore imperative that sampling is guided.

One way to guide the sampling procedure is with the use of patch level annotations. For example, on breast cancer metastasis detection, in multiple papers, patches from normal and tumor regions were extracted based on pixel level labels that were provided by pathologists (14, 15, 23, 24, 36–38). Others were able



to detect, segment, and classify different types of cell nuclei, colon glands and other large organs, as well as to classify and localize a variety of diseases (35, 39, 40). Most successful approaches also employ hard negative mining, an iterative process whereby false positives are added to the training data set for further training (14, 23, 24, 36). Identification of false positives is possible in such cases due to the availability of patch level annotations.

### 3.1.2. Slide Level Annotation

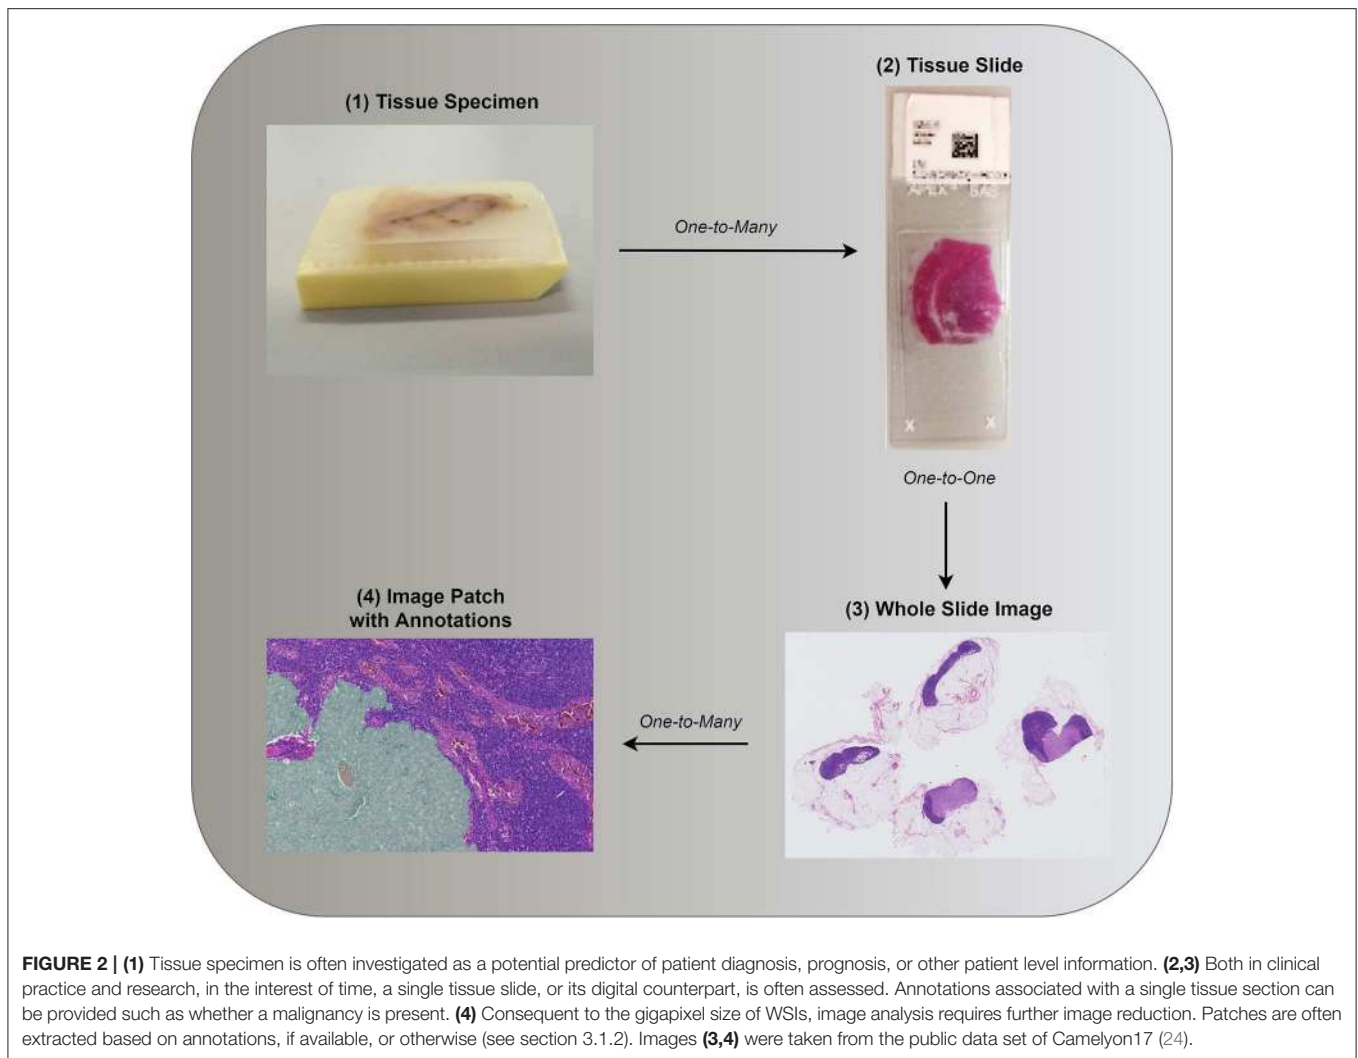
Due to practical limitations, in most cases ground truth labeling is done on the level of WSIs as opposed to individual patches. Despite this lower granularity of labeling, a number of deep learning based approaches have demonstrated highly promising results. Techniques vary and often take on the form of multiple instance learning, unsupervised learning, reinforcement learning, and transfer learning, or a combination of thereof. Intuitively, the goal is usually to identify patches that can collectively or independently predict the whole slide label.

Preprocessing based on image filters can be employed to reduce the number of patches that need to be analyzed. Multiple studies also employ the Otsu, hysteresis, or other types of threshold as an automatic way of identifying tissue within the WSI. Other operations such as contrast normalization, morphological operations, and a problem specific patch scoring system can also be employed to reduce further the number of candidate patches and even enable automatic patch localization.

However, verifying that indeed each patch has the same label as the slide often requires domain-specific expertise and even the process of coming up with the best image filters requires at the very least some human intuition.

In order to avoid potential human bias, most approaches employ unsupervised or multi-instance learning, or a combination of both. Tellez et al. (18) examined different methods of unsupervised representation learning for compressing image patches into a lower dimensional latent space. It was then possible to train CNNs directly on the compressed WSIs (18). Others reduced the dimensionality of patches using traditional dimensionality reduction techniques, such as principal component analysis, as well as CNNs pretrained on ImageNet (16, 41). Both Zhu et al. (41) and Yue et al. (16) subsequently used *k*-means clustering and found the most discriminative clusters of patches by training CNNs in a weakly supervised manner (42). However, the multi-stage structure of the aforementioned techniques does not allow processes that come first, i.e., patch compression or patch localization, to improve following the improvement of later processes, i.e., visual understanding.

Several other ideas enable persistent improvement of patch localization and visual understanding either by iteratively revising each process or by learning both in an end-to-end fashion. Hou et al. (43) proposed an expectation maximization algorithm that enables increasingly more discriminative patches



to be selected while at each iteration the CNN was trained further for 2 epochs. A patch was considered more discriminative if, when given as an input to the CNN, the prediction was closer to the slide level label (43). Combalia and Vilaplana (44) instead of an expectation maximization algorithm, employed Monte Carlo sampling. One of the most promising emerging directions aims at incorporating the process of selecting patches within the optimization of visual understanding (17, 45–49). Courtiol et al. (45) modified a pretrained ResNet-50 by adding a  $1 \times 1$  convolutional layer after the convolutional layers to get patch level predictions. A MinMax layer was added on top followed by fully connected layers to predict the slide level label (45). The MinMax layer is a type of attention mechanism which gives the capability of selective training on the most discriminative patches of both classes.

Instead of extracting features from all or most patches before selecting a few to learn on, recent work has employed attention models (17, 47–49). Intuitively, an attention model is initially as good as random guessing at patch selection but progressively chooses more discriminative patches that contribute to better model performance. For example, Qaiser and Rajpoot (47) used

reinforcement learning to train a model in selecting patches at  $20\times$  and  $10\times$  magnification levels based on a low resolution image at  $2.5\times$  magnification level. Using supervised learning, BenTaieb and Ghassan Hamarneh (17) employed a recurrent visual attention network that processes non-overlapping patches of  $5,000 \times 5,000$  pixels at  $20\times$  magnification level and sequentially identifies and analyses different regions within those patches.

### 3.1.3. Patient Level Annotation

Usually, multiple WSIs can be acquired for each patient since the initial tissue occupies a 3D space and therefore multiple cuts can be made. In this case, the available ground truth can be specific to the patient, but not to each individual WSI (16, 41). This is typically addressed based on the same operations that are used when aggregating patch level predictions to slide level predictions (16, 41).

### 3.1.4. Aggregating to a Higher Level

In many cases, training takes place at a lower level, e.g., patch level, but the end goal resides at a higher level, e.g., slide level.

For example, in the case of cancer diagnosis, a CNN may be trained to identify the presence of cancerous cells within a patch. However, some type of aggregation is needed in order to infer whether a WSI contains cancerous cells. This may take the form of a maximum or average operation over some or all patch predictions. In other cases, traditional machine learning models, or recurrent neural networks may be employed and trained using features extracted by a CNN and the ground truth that is available at a higher level.

### 3.2. Beyond Patch Extraction

A primary limitation of patch based analysis emerges as a consequence of analysing a large input image by means of independent analysis of smaller regions. In particular, such approaches are inherently unable to capture information distributed over scales greater than the patch size. For example, although cell characteristics can be extracted from individual patches, higher level structural information, such as the shape or extend of a tumor, can only be captured when analysing larger regions. Explicitly modeling spatial correlations between patches has been proposed as a potential solution (36–38). However, this idea has only been tested with a small number of neighborhoods and requires patch level annotations. A different approach involves patch extraction from multiple magnification levels (15, 27, 43). Others, such as the attention models described above, consider global context, that is a low resolution image of the WSI, both when choosing regions to attend to and when predicting the slide level label (17, 47). Finally, recent work attempts to ameliorate some of the aforementioned problems associated with patch based analysis by using much larger patch sizes (17, 48).

## 4. DISCUSSION

The aim of computer vision is to create algorithmic solutions capable of visual understanding. Applications can range from object identification and detection to image captioning and scene decomposition. In the past decade most areas of computer vision have seen remarkable progress, much of it effected by advances in neural network based learning algorithms (50, 51). The success of these methodologies, part of the now established field of deep learning, can be attributed to a number of reasons, with the transformation of the feature extraction stage often described as the leading factor.

In the previous decade most approaches focused on finding ways to explicitly extract features from images for models subsequently to employ (52, 53). Therefore, feature extraction and model development were two distinct, independent stages that were performed sequentially, and where the former was based on human intuition of what constitutes a good feature. Automating this process through the use of convolutional neural networks (CNNs) has been shown to result in more discriminative features tailored for the problem at hand (13, 23, 25, 26). This is one of the reasons behind the success of deep learning, and more broadly, neural network based learning, as feature extraction became a learning process, fundamentally

intertwined with the learning of model parameters. Some of the other key factors which contributed to the successes of deep learning include advancements in hardware and software, as well as the increase in data availability.

The analysis of multi-gigabyte images is a new challenge for deep learning that has only appeared along the emergence of digital pathology and whole slide imaging. Building deep learning models capable of understanding WSIs presents novel challenges to the field. When patch level labels are available, patch sampling coupled with hard negative mining can train deep learning models that in many cases match and even surpass the accuracy of pathologists (13). For many medical data sets with patch level annotations, deep learning models seem to excel, and with the introduction of competitions, such as Camelyon16 and Camelyon17 (23, 24), this type of deep learning has repeatedly demonstrated its success in performance and interpretability (6). Therefore, patch based learning from gigapixel images and patch level annotations seems to be the closest to clinical employment. However, in many cases, only labeling with lower granularity can be attained either because it is very laborious and expensive, or simply because it is infeasible. In addition, patch level supervision may be limiting the potential of deep learning models as the models can only be as good as the annotations provided.

To work with slide or patient level labels, current approaches focus on the *where* problem, or in other words, on approximating the spatial distribution of the signal (18). Out of the work we reviewed, only Tellez et al. (18) has instead simplified the *what* problem, i.e., visual understanding, to the point where the *where* problem becomes trivial. It would be interesting to see the efficacy of the work by Tellez et al. on harder problems, such as prognosis estimation, and with other low dimensional latent mappings. On the *where* problem, there are generally two approaches. The first uses a type of meta-learning, where in order to optimize the *where* problem, the *what* problem has to first be optimized. The second approach attempts to optimize both *what* and *where* problems simultaneously in an end-to-end setting. This is done by either forwarding a set of patches through a CNN and attending on a few or by localizing and attending to a single patch at each time step.

Deep learning is already demonstrating its potential across a wide range of medical problems associated with digital pathology. However, the need for detailed annotations limits the applicability of strongly supervised techniques. Other techniques from weakly supervised, unsupervised, reinforcement, and transfer learning are employed to counter the need for detailed annotations while dealing with massive, highly heterogeneous images and small data sets. This emerging direction away from strong supervision opens new opportunities in WSI analysis, such as addressing problems for which the ground truth is only known at a higher than patch level, e.g., patient survivability and recurrence prediction.

## AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

## REFERENCES

- Snead RJD, Tsang Y, Meskiri A, Kimani KP, Crossman R, Rajpoot MN, et al. Validation of digital pathology imaging for primary histopathological diagnosis. *Histopathology*. (2015) 68:1063–72. doi: 10.1111/his.12879
- Pantanowitz L, Sinard HJ, Henricks HW, Fatheree AL, Carter BA, Contis L, et al. Validating whole slide imaging for diagnostic purposes in pathology: guideline from the college of American pathologists pathology and laboratory quality center. *Arch Pathol Lab Med*. (2013) 137:1710–22. doi: 10.5858/arpa.2013-0093-CP
- Cacomo S. *FDA Allows Marketing of First Whole Slide Imaging System for Digital Pathology* (2017). Available online at: <https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm552742.htm>
- Hamilton WP, Bankhead P, Wang Y, Hutchinson R, Kieran D, McArt GD, et al. Digital pathology and image analysis in tissue biomarker research. *Methods*. (2014) 70:59–73. doi: 10.1016/j.ymeth.2014.06.015
- Harder N, Schönmeier R, Nekolla K, Meier A, Brieu N, Vanegas C, et al. Automatic discovery of image-based signatures for ipilimumab response prediction in malignant melanoma. *Sci Rep*. (2019) 9:7449. doi: 10.1038/s41598-019-43525-8
- Janowczyk A, Madabhushi A. Deep learning for digital pathology image analysis: a comprehensive tutorial with selected use cases. *J Pathol Informat*. (2016) 7:29. doi: 10.4103/2153-3539.186902
- Brieu N, Gavriel GC, Nearchou PI, Harrison JD, Schmidt G, Caie DP. Automated tumour budding quantification by machine learning augments TNM staging in muscle-invasive bladder cancer prognosis. *Sci Rep*. (2019) 9:5174. doi: 10.1038/s41598-019-41595-2
- Caie DP, Zhou Y, Turnbull KA, Oniscu A, Harrison JD. Novel histopathologic feature identified through image analysis augments stage II colorectal cancer clinical reporting. *Oncotarget*. (2016) 7:44381–94. doi: 10.18632/oncotarget.10053
- Beck HA, Sangoi RA, Leung S, Marinelli JR, Nielsen OT, van de Vijver JM, et al. Systematic analysis of breast cancer morphology uncovers stromal features associated with survival. *Sci Transl Med*. (2011) 3:108ra113. doi: 10.1126/scitranslmed.3002564
- Nearchou PI, Lillard K, Gavriel GC, Ueno H, Harrison JD, Caie DP. Automated analysis of lymphocytic infiltration, tumor budding, and their spatial relationship improves prognostic accuracy in colorectal cancer. *Cancer Immunol Res*. (2019) 7:609–20. doi: 10.1158/2326-6066.CIR-18-0377
- Sari TC, Gunduz-Demir C. Unsupervised feature extraction via deep learning for histopathological classification of colon tissue images. *IEEE Trans Med Imaging*. (2019) 38:1139–49. doi: 10.1109/TMI.2018.2879369
- Dimitriou N, Arandjelović O, Harrison JD, Caie DP. A principled machine learning framework improves accuracy of stage II colorectal cancer prognosis. *npj Digit Med*. (2018) 1:52. doi: 10.1038/s41746-018-0057-x
- Qaiser T, Mukherjee A, Reddy PB C, Munugoti DS, Tallam V, Pitkäaho T, et al. HER2 challenge contest: a detailed assessment of automated HER2 scoring algorithms in whole slide images of breast cancer tissues. *Histopathology*. (2017) 72:227–38. doi: 10.1111/his.13333
- Wang D, Khosla A, Gargeya R, Irshad H, Beck HA. Deep learning for identifying metastatic breast cancer. *arXiv*. (2016).
- Liu Y, Gadepalli K, Norouzi M, Dahl EG, Kohlberger T, Boyko A, et al. Detecting cancer metastases on gigapixel pathology images. *arXiv*. (2017).
- Yue X, Dimitriou N, Caie DP, Harrison JD, Arandjelovic O. Colorectal cancer outcome prediction from H&E whole slide images using machine learning and automatically inferred phenotype profiles. In: *Conference on Bioinformatics and Computational Biology*. Vol. 60, Honolulu, HI (2019). p. 139–49.
- BenTaieb A, Ghassan Hamarneh G. Predicting cancer with a recurrent visual attention model for histopathology images. In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G, editors. *Medical Image Computing and Computer-Assisted Intervention*. Vol. 11071. Springer International Publishing (2018). p. 129–37. Available online at: <https://link.springer.com/book/10.1007/978-3-030-00934-2>
- Tellez D, Litjens G, van der Laak J, Ciompi F. Neural image compression for gigapixel histopathology image analysis. *arXiv*. (2018). doi: 10.1109/TPAMI.2019.2936841
- Wong FP, Wei W, Smithy WJ, Acs B, Toki IM, Blenman RMK, et al. Multiplex quantitative analysis of tumor-infiltrating lymphocytes and immunotherapy outcome in metastatic melanoma. *Clin Cancer Res*. (2019) 25:2442–9. doi: 10.1158/1078-0432.CCR-18-2652
- Schlag I, Arandjelović O. Ancient Roman coin recognition in the wild using deep learning based recognition of artistically depicted face profiles. In: *Proceedings of the International Conference on Computer Vision*. Venice (2017). p. 2898–906.
- Cooper J, Arandjelović O. Visually understanding rather than merely matching ancient coin images. In: *Proceedings of the INNS Conference on Big Data and Deep Learning*. Sestri Levante (2019).
- Sirinukunwattana K, Domingo E, Richman S, Redmond LK, Blake A, Verrill C, et al. Image-based consensus molecular subtype classification (imCMS) of colorectal cancer using deep learning. *bioRxiv*. (2019). doi: 10.1101/645143
- Bejnordi EB, Veta M, van Diest JP, van Ginneken B, Karssemeijer N, Litjens G, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *J Am Med Assoc*. (2017) 318:2199–210. doi: 10.1001/jama.2017.14580
- Bandi P, Geessink O, Manson Q, van Dijk M, Balkenhol M, Hermsen M, et al. From detection of individual metastases to classification of lymph node status at the patient level: the CAMELYON17 challenge. *IEEE Trans Med Imaging*. (2019) 38:550–560. doi: 10.1109/TMI.2018.2867350
- Aresta G, Araújo T, Kwok S, Chennamsetty SS, Safwan M, Alex V, et al. BACH: grand challenge on breast cancer histology images. *arXiv*. (2018).
- Veta M, Heng JY, Stathonikos N, Bejnordi EB, Beca F, Wollmann T, et al. Predicting breast tumor proliferation from whole-slide images: the TUPAC16 challenge. *Med Image Anal*. (2019) 54:111–21. doi: 10.1016/j.media.2019.02.012
- Campanella G, Silva WKV, Fuchs JT. Terabyte-scale deep multiple instance learning for classification and localization in pathology. *arXiv*. (2018).
- Bennett W, Smith K, Jarosz Q, Nolan T, Bosch W. Reengineering workflow for curation of DICOM datasets. *J Digit Imaging*. (2018) 31:783–91. doi: 10.1007/s10278-018-0097-4
- Kahn EC, Carrino AJ, Flynn JM, Peck JD, Horii CS. DICOM and radiology: past, present, and future. *J Am Coll Radiol*. (2007) 4:652–7. doi: 10.1016/j.jacr.2007.06.004
- Lennerz KJ, Herrmann DM, Clunie AD, Fedorov A, Doyle WS, Pieper S, et al. Implementing the DICOM standard for digital pathology. *J Pathol Informat*. (2018) 9:37. doi: 10.4103/jpi.jpi\_42\_18
- Clunie AD. Dual-personality DICOM-TIFF for whole slide images: a migration technique for legacy software. *J Pathol Informat*. (2019) 10:12. doi: 10.4103/jpi.jpi\_93\_18
- Caie DP, Schuur K, Oniscu A, Mullen P, Reynolds AP, Harrison JD. Human tissue in systems medicine. *FEBS J*. (2013) 280:5949–56. doi: 10.1111/febs.12550
- Magee RD, Treanor D, Crellin D, Shires M, Smith EK, Mohee K, et al. Colour normalisation in digital histopathology images. In: *Proc. Opt. Tissue Image Anal. Microsc. Histopathol. Endosc*. London (2009). p. 100–11.
- Brieu N, Caie DP, Gavriel GC, Schmidt G, Harrison JD. Context-based interpolation of coarse deep learning prediction maps for the segmentation of fine structures in immunofluorescence images. In: *Medical Imaging 2018: Digital Pathology*. Houston, TX (2018).
- Chang YH, Jung KC, Woo IJ, Lee S, Cho J, Kim WS, et al. Artificial intelligence in pathology. *J Pathol Transl Med*. (2019) 53:1–12. doi: 10.4132/jptm.2018.12.16
- Li Y, Ping W. Cancer metastasis detection with neural conditional random field. *arXiv*. (2018).
- Kong B, Wang X, Li Z, Song Q, Zhang S. Cancer metastasis detection via spatially structured deep network. In: *Information Processing in Medical Imaging*. Boone, NC: Springer International Publishing (2017). p. 236–48.
- Zanjani GF, Zinger S, de With HNP. Cancer detection in histopathology whole-slide images using conditional random fields on deep embedded spaces. In: *Medical Imaging*. Vol. 10581, Houston, TX: SPIE (2018).
- Litjens G, Kooi T, Bejnordi EB, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. *Med Image Anal*. (2017) 42:60–88. doi: 10.1016/j.media.2017.07.005

40. Coudray N, Ocampo SP, Sakellaropoulos T, Narula N, Snuderl M, Fenyö D, et al. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat Med.* (2018) 24:1559–67. doi: 10.1038/s41591-018-0177-5
41. Zhu X, Yao J, Zhu F, Huang J. WSISA: making survival prediction from whole slide histopathological images. In: *IEEE Conference on Computer Vision and Pattern Recognition*. Honolulu, HI (2017). p. 6855–63. doi: 10.1109/CVPR.2017.725
42. Zhou Z. A brief introduction to weakly supervised learning. *Natl Sci Rev.* (2017) 5:44–53. doi: 10.1093/nsr/nwx106
43. Hou L, Samaras D, Kurc MT, Gao Y, Davis EJ, Saltz HJ. Patch-based convolutional neural network for whole slide tissue image classification. In: *IEEE Conference on Computer Vision and Pattern Recognition*. Las Vegas, NV (2016). p. 2424–33. doi: 10.1109/CVPR.2016.266
44. Combalia M, Vilaplana V. Monte-Carlo sampling applied to multiple instance learning for histological image classification. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Granada (2018). p. 274–81.
45. Courtiol P, Tramel WE, Sanselme M, Wainrib G. Classification and disease localization in histopathology using only global labels: a weakly-supervised approach. *arXiv.* (2018).
46. Ilse M, Tomczak MJ, Welling M. Attention-based deep multiple instance learning. *arXiv.* (2018).
47. Qaiser T, Rajpoot MN. Learning where to see: a novel attention model for automated immunohistochemical scoring. *arXiv* (2019). doi: 10.1109/TMI.2019.2907049
48. Momeni A, Thibault M, Gevaert O. Deep recurrent attention models for histopathological image analysis. *bioRxiv.* (2018). doi: 10.1101/438341
49. Tomita N, Abdollahi B, Wei J, Ren B, Suriawinata A, Hassanpour S. Finding a needle in the haystack: attention-based classification of high resolution microscopy images. *arXiv.* (2018).
50. Guo Y, Liu Y, Oerlemans A, Lao S, Wu S, Lew SM. Deep learning for visual understanding: a review. *Neurocomputing.* (2016) 187:27–48. doi: 10.1016/j.neucom.2015.09.116
51. Voulodimos A, Doulamis N, Doulamis A, Protopapadakis E. Deep learning for computer vision: a brief review. *Comput Intell Neurosci.* (2018) 2018:1–13. doi: 10.1155/2018/7068349
52. Lowe GD. Distinctive image features from scale-invariant keypoints. *Int J Comput Vis.* (2004) 60:91–110. doi: 10.1023/B:VISI.0000029664.99615.94
53. Dalal N, Triggs B. Histograms of oriented gradients for human detection. In: *IEEE Conference on Computer Vision and Pattern Recognition*. Vol. 1, San Diego, CA (2005). p. 886–93.

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Dimitriou, Arandjelović and Caie. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.